

# Leveraging Activation Maximization and Generative Adversarial Training to Recognize and Explain Patterns in Natural Areas in Satellite Imagery

Ahmed Emam<sup>id</sup>, Timo T. Stomberg<sup>id</sup>, and Ribana Roscher<sup>id</sup>, *Member, IEEE*

**Abstract**—Natural protected areas are vital for biodiversity, climate change mitigation, and supporting ecological processes. Despite their significance, comprehensive mapping is hindered by a lack of understanding of their characteristics and a missing land-cover class definition. This letter aims to advance the explanation of the designating patterns forming protected and wild areas. To this end, we propose a novel framework that uses activation maximization and a generative adversarial model. With this, we aim to generate satellite images that, in combination with domain knowledge, are capable of offering complete and valid explanations for the spatial and spectral patterns that define the natural authenticity of these regions. Our proposed framework produces more precise attribution maps pinpointing the designating patterns forming the natural authenticity of protected areas. Our approach fosters our understanding of the ecological integrity of the protected natural areas and may contribute to future monitoring and preservation efforts.

**Index Terms**—Activation maximization, explainable machine learning, generative models, patterns discovery.

## I. INTRODUCTION

NATURAL areas with minimum human influence refer to regions on Earth that remain largely untouched by human activities, such as urbanization, agriculture, logging, and other forms of development. These areas are characterized by high levels of biodiversity and provide a wealth of ecological integrity benefits; for example, these areas offer unique access to understanding nearly undisrupted natural ecosystem processes, such as water and pollination cycles. Therefore, comprehensive mapping and monitoring of highly protected natural areas are crucial to understanding the geocological patterns essential for these areas to thrive [1]. With this in mind, it is not surprising that the monitoring and understanding of natural areas have gained significant attention in recent years within the remote sensing and ecological research communities.

Manuscript received 24 May 2023; revised 21 September 2023 and 14 November 2023; accepted 19 November 2023. Date of publication 21 November 2023; date of current version 19 December 2023. This work was supported in part by the Deutsche Forschungsgemeinschaft [German Research Foundation (DFG)] under Germany's Excellence Strategy, under Grant EXC-2070-390732324-PhenoRob. (*Corresponding author: Ahmed Emam.*)

Ahmed Emam and Timo T. Stomberg are with the Institute of Geodesy and Geoinformation, University of Bonn, 53113 Bonn, Germany (e-mail: aemam@uni-bonn.de; timo.stomberg@uni-bonn.de).

Ribana Roscher is with the Institute of Bio- and Geosciences, Forschungszentrum Jülich, 52428 Jülich, Germany, and also with the Institute of Geodesy and Geoinformation, University of Bonn, 53113 Bonn, Germany (e-mail: r.roscher@fz-juelich.de).

Digital Object Identifier 10.1109/LGRS.2023.3335473

Satellite imagery presents an effective approach for continuously monitoring vast protected natural areas that are challenging for humans to access. By utilizing this technology, efficient and cost-effective data collection becomes possible while minimizing disturbances to sensitive ecosystems. Leveraging large satellite imagery datasets, machine learning models, such as deep convolutional neural networks (CNNs), can accurately differentiate natural areas. For instance, in a study conducted by Ekim et al. [2], a dataset and a baseline CNN model were developed to precisely identify and classify protected natural areas.

Although these models can effectively detect specific patterns that characterize such areas, the patterns are mostly not inheritably accessible or interpretable by humans. Therefore, explainable machine learning techniques are utilized to explain the designating patterns that drive the model's decision-making process. In the context of analyzing natural areas, Stomberg et al. [3] used an inheritable explainable classification network that produces attribution maps that localize patterns characterizing protected natural areas in satellite imagery. Standard explainable machine learning techniques, such as occlusion sensitivity maps [4], gradient-based class activation maps [5], and deep learning important features (DeepLIFT) [6], can identify influential pixels within the image and their significance in deriving the machine learning model's decision. While these methods can provide partial explanations, they lack the capacity to generate complete and valid explanations that accurately characterize natural areas with minimal human interference. For attributions to be considered valid and complete, they must assign high values to the class-discriminative patterns and be consistent with expert opinions [7], [8].

Our contribution introduces a novel approach that leverages activation maximization [9] and draws inspiration from cycle consistent generative adversarial networks (Cycle GANs) [10]. Our framework, when combined with domain knowledge, offers complete and valid explanations for the patterns that define natural areas with modern human impact. Activation maximization is an explainable machine learning technique that modifies an input image to maximize its classification score produced by CNN. By extending this idea, we utilize activation maximization in our proposed modified Cycle GAN's objective function to generate two images that maximize and minimize the classification score produced by a CNN. By optimizing the Cycle GAN for maximizing and minimizing the classification score, the two generated

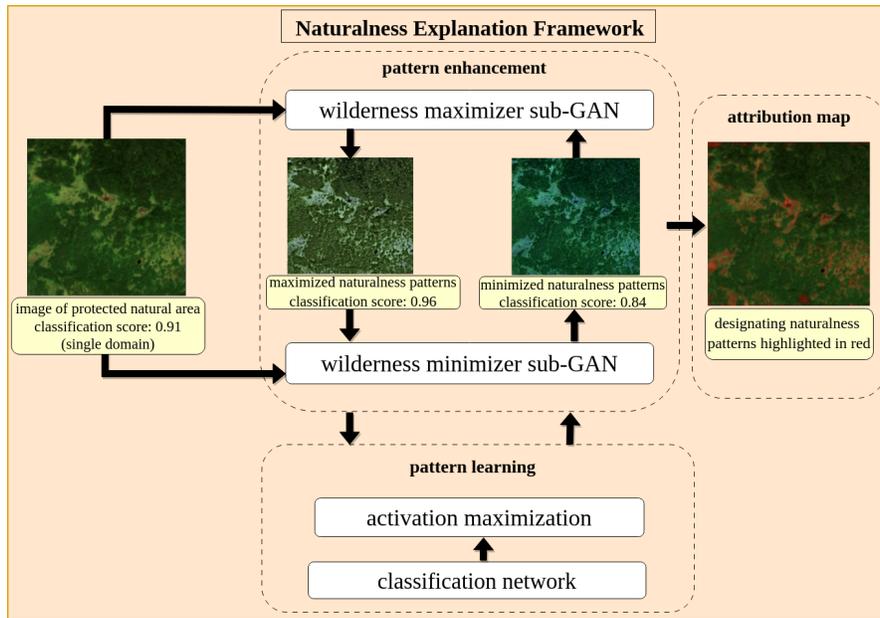


Fig. 1. Flowchart of the proposed approach. (Left) Input image displays the class of interest, and it has a classification score of 0.91 from the classification network. (Top middle) Pattern enhancement—Cycle GAN consists of the pattern-maximizer sub-GAN that generates an image with enhanced patterns of the specific class and the pattern-minimizer sub-GAN that generates an image of the same size but with suppressed characteristics. The synthesized images have the classification scores of 0.96 and 0.84, respectively. (Bottom middle) Pattern learning—trained classifier for the class of interest being fed with the generated images from the pattern enhancement part, giving feedback to both generators in the Cycle GAN [10] to maximize and minimize the characteristics in both sub-GANs, respectively. (Right) Attribution mapping—showing the contribution of each pixel to the class of interest. To generate the attribution map, we subtract each pixel value in the generated image from the input image and calculate the average difference over all channels. We overlaid the original image with the average absolute difference. We have designed a novel framework that modifies the Cycle GAN’s objective to generate both maximized and minimized naturalness patterns, enabling us to create more precise attribution maps. To the best of our knowledge, we are pioneering this framework for explaining naturalness patterns in protected areas.

images will exhibit slight variations in their color and texture. By comparing these images, an attribution map with complete and valid explanations can be derived that depicts the characteristic patterns of protected natural areas with minimal human interference. Katzmann et al. [11] used a comparable approach for explaining medical decisions made by machine learning models. To the best of our knowledge, this work is the first that integrates activation maximization in the Cycle GAN objective function and incorporates domain knowledge to interpret designating patterns contributing to the concept of naturalness in protected natural areas.

## II. PROPOSED APPROACH

Our approach involves three consecutive phases (see Fig. 1) and is used to analyze and interpret designating patterns of protected natural areas.

### A. Pattern Learning

In the pattern learning phase, we train a single-class classifier to learn the characteristic and distinctive patterns of a class of interest. To enable this, we use a network that has a single output, which can be defined either in the style of a regression network as a score in the range [0, 1], as proposed in [3] or as a binary decision classifier.

### B. Pattern Enhancement

The pattern enhancement phase is conducted after completing the training process of the classifier in the pattern learning phase. Here, we integrate the concept of activation

maximization in our modified cyclic consistency GAN’s objective function to generate two images with pattern-maximized and pattern-minimized image characteristics. In contrast to commonly used GANs, as the ones proposed by Zhu et al. [10] and Katzmann et al. [11] that are primarily designed for performing domain transfer from one image domain to another image domain, our approach has only a single input domain, and the generated outputs would be of the same domain as the input’s domain with minimum changes pinpointing the important patterns in the input image influencing the classifier’s decisions.

1) *Architecture*: We modify the original Cycle GAN’s architecture and the objective function, as presented in [10], to generate two images based on the input image with minimum changes while keeping the input image’s main structures unchanged. These changes lead to the maximization or minimization of the classification score of the generated images. Our Cycle GAN consists of two sub-GANs, namely, the pattern maximizer and the pattern minimizer. Each sub-GAN consists of two neural networks, a generator  $w$  and a discriminator  $d$ . Both generators are ResNet-based CNNs [12] with 12 residual blocks. Each generator produces two modified versions of the input image of the same size as the input image. The generator’s architecture is illustrated in Fig. 2.

2) *Loss Function*: The discriminators in both sub-GANs are trained using binary cross entropy

$$\mathcal{L}_{CE}(h, \hat{h}) = h \cdot \log(\hat{h}) + (1 - h) \cdot \log(1 - \hat{h}). \quad (1)$$

The scalar  $h$  is a binary variable indicating whether a sample is real or generated, and  $\hat{h}$  is the discriminator’s prediction.

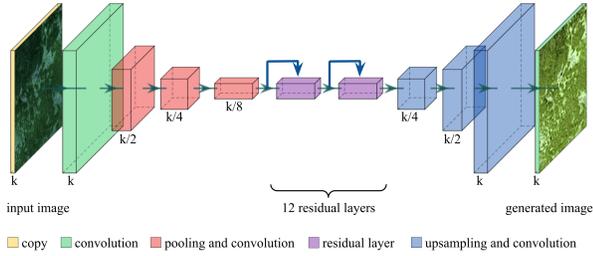


Fig. 2. Generator's architecture used in both sub-GANs in the pattern enhancement phase.

The overall generator's loss function in both sub-GANs consists of four terms. The similarity loss term enforces the similarity between the generated images and the original input image

$$\mathcal{L}_{\text{sim}}(X, \hat{X}) = \left| X - \hat{X} \right|_F \quad (2)$$

where  $X$  is the input image and  $\hat{X}$  is the generated pattern-maximized or pattern-minimized output image. The cycle consistency loss term enforces that successive mapping through both generators leads to a reconstruction close to the original input image

$$\mathcal{L}_{\text{cyc}}(w^-, w^+, X) = \left| X - (w^+(w^-(X))) \right| \quad (3)$$

where  $w^+$  is the pattern-maximizing sub-GAN generator and  $w^-$  is the pattern-minimizing sub-GAN generator. Moreover, the cycle consistency loss enforces the similarity between the two generated images and the input image. The adversarial training loss guides the generator  $w$  to generate realistic images following the guidance of the discriminator  $d$ :

$$\mathcal{L}_{\text{adv}}(d, w, X) = (d(w(X)) - 1)^2. \quad (4)$$

The activation maximization loss term, as proposed in [11] and [13], compels the generator to add features to the generated images that will increase and decrease its classification score for both the pattern-maximizer and pattern-minimizer generators, respectively,

$$\mathcal{L}_{\text{AM}}(y, w, X) = \mathcal{L}_{\text{CE}}(y, w(X)) \quad (5)$$

where  $y$  is the label, which is 1 for the pattern maximizer and 0 for the pattern minimizer, with  $\mathcal{L}_{\text{CE}}$  computed using (1). For both similarity and cycle consistency loss terms, we use the  $L_1$  loss function. We use the binary cross-entropy loss function for the adversarial and activation maximization loss.

The training of both sub-GANs is performed in two steps. In the pretraining phase, we use the loss term without activation maximization loss term to initialize the weights of both sub-GANs and to ensure that both generators can reconstruct the input image. The main training was performed using the complete generator loss function, which is defined as follows for the pattern maximizer:

$$\begin{aligned} \mathcal{L}_{\text{max}}(X, y, w^+, w^-, d^+) &= \mathcal{L}_{\text{sim}}(X, \hat{X}) \\ &+ \mathcal{L}_{\text{cyc}}(w^-, w^+, X) \mathcal{L}_{\text{adv}}(d^+, w^+, X) \\ &+ \lambda \mathcal{L}_{\text{AM}}(y, w^+, X) \end{aligned} \quad (6)$$

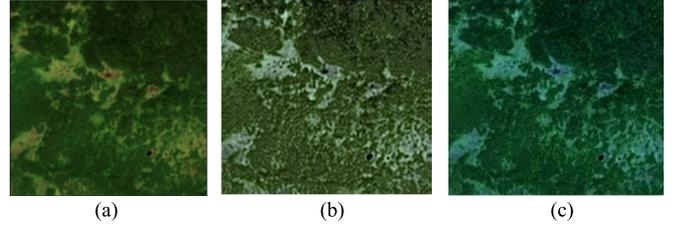


Fig. 3. Generated outputs from the sub-GANs. (a) Original image. (b) Output of the pattern maximizer sub-GAN. (c) Output of the pattern minimizer sub-GAN.

where  $\lambda$  is the weighting factor assigned to the activation maximization loss term,  $w^-$  being the pattern-minimizing generator, and  $d^-$  its respective discriminator. The loss function is applied accordingly to the pattern-minimizer generator  $w^-$  by inverting  $(\cdot)^+$  and  $(\cdot)^-$  and setting  $y = 0$ .

### C. Attribution Mapping

To produce the attribution maps of the characterizing patterns, we calculate the average absolute difference between the generated images for each image band. Higher attributions are assigned to the areas with a higher average absolute difference, because these areas are related to the designating patterns of the class of interest.

## III. EXPERIMENTS, RESULTS, AND DISCUSSION

In our use case, we aim to deepen our understanding of the most influential characteristics in the protected natural areas in Fennoscandia (Norway, Sweden, and Finland) by interpreting and explaining the designating patterns that contribute to the naturalness authenticity of this region.

### A. Dataset and Experimental Setup

1) *AnthroProtect Dataset*: We use the AnthroProtect dataset [3] consisting of approximately 24 000 multispectral Sentinel-2 images of size  $256 \times 256$  pixels, showing protected or anthropogenic areas in Fennoscandia. The protected natural areas are chosen based on the World Database on Protected Areas (WDPA) [14]. The protected natural areas are located within the categories “strict nature reserve” (Ia), “wilderness” (Ib), and “national park” (II). The anthropogenic areas are within or close to areas dominated by the CORINE land-cover classes “artificial surfaces” and “agricultural areas,” while the CORINE dataset represents that domain knowledge [15]. The AnthroProtect dataset, thus, contains images that lie somewhere at the extremes of the unspecific, continuous wilderness scale within Fennoscandia. To minimize the computational cost, we restrict ourselves to using the red, green, and blue bands; however, our proposed framework is flexible regarding the dimensionality of the input. The AnthroProtect dataset has been published with a proposed data split (training, validation, and test), which we adopt accordingly.

In the pattern learning phase, any classification neural network can be employed to learn the designating patterns of naturalness-protected areas. We opted for a modified version of the classification network from [3], with only three input

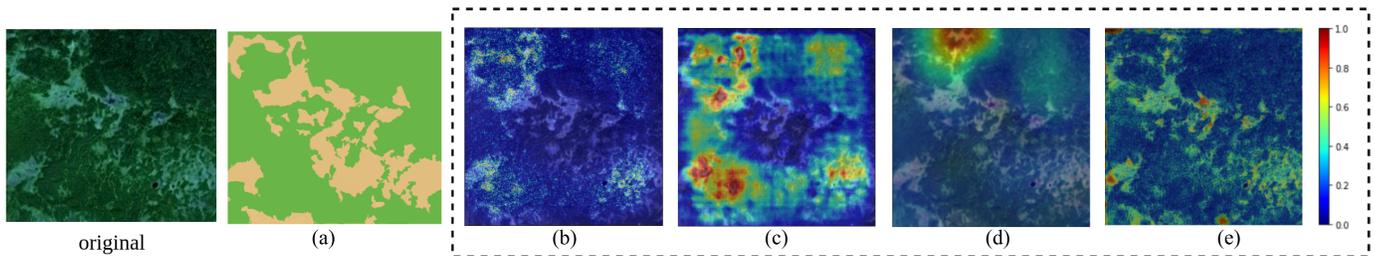


Fig. 4. Comparison of attribution maps from different explainable machine learning methods. (a) Segmentation mask showing the wetlands pattern in yellow and green representing heavy vegetation. (b) Attribution map from DeepLIFT. (c) Occlusion sensitivity map. (d) GradCAM. (e) Attribution map from our proposed method. The scale describes plots in the dashed rectangle and assesses the contribution of pixels to the authenticity of protected natural areas. A score of 1 signifies the highest possible contribution. A score of 0 signifies no contribution to the naturalness of protected areas.

channels instead of the ten channels as the network in [3] to conserve computational resources. The use of fewer input channels led to a minor decrease in the classification accuracy but did not have a negative effect on the performance of the pattern enhancement phase. The training, validation, and test set accuracies for the classification network are 98.7%, 97.3%, and 99.3%, respectively.

During the network's training phase, we adopted the one-cycle learning schedule strategy [16] with a maximum learning rate set at 0.01. To optimize the model's parameters, we employed gradient descent, incorporating a weight decay of 0.0001. Our training data used a batch size of 32, and to ensure consistency, the pixel values of Sentinel-2 images were normalized to a range between 0 and 1 by dividing them by 10 000.

We incorporated CutMix [17] during training. This technique allowed us to create synthetic images by mixing portions of areas with human influence and regions devoid of anthropogenic impact. This augmentation aids the classifier in evaluating the naturalness integrity of an image, with a minimum value of 0 representing full anthropogenic influence and a maximum value of 1 denoting completely untouched natural areas.

### B. Our Modified Cycle GAN

We use the PatchGAN [18] for both discriminators. Their goal is to classify overlapping  $70 \times 70$  image patches as either generated or real and to provide feedback to the generator to produce a realistic output. This architecture pushes the discriminator not to focus solely on the artifacts in the generated image but on the general quality of the generated image, which leads to beneficiary feedback to the generator during the Cycle GAN's training process [10]. We set  $\lambda_{AM} = 0.3$ , resulting in high-resolution outputs with minimum artifacts. We used the Adam optimizer [19] to train the generative model with the decay of first-order momentum of gradients  $\beta_1 = 0.5$  and  $\beta_2 = 0.999$ .

### C. Results and Discussion

Figs. 1 and 3 show the example images generated by the pattern-maximizer and the pattern-minimizer sub-GANs. The pattern maximizer produces images with higher pixel values in the green channel compared with the original image, whereas the pattern minimizer produces the same effect on the blue channel. Including environmental domain knowledge

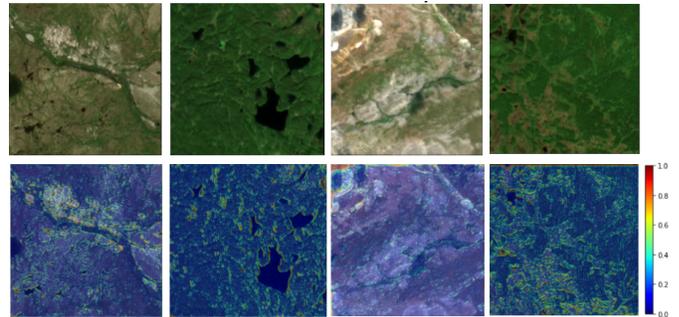


Fig. 5. (Top) Original input images. (Bottom) Corresponding attribution maps produced by our approach.

and information from the existing land-cover products, we can see that brighter wetlands areas were produced in the pattern-maximized images, and the borders of wetlands and the surroundings of water bodies are highlighted with vibrant green color by the pattern minimizer.

We also depict attribution maps showing the importance of each pixel in the input image, contributing to the naturalness integrity and its influence on the classifier's decision; see Fig. 4. By incorporating our method's attribution map with the correspondent segmentation mask from the CORINE dataset, we can conclude that our method assigns high importance to wetlands and the small water bodies (bonds) in the middle of the image and lower attributions to the nondiscriminative vegetation areas between the characterizing features. The occlusion sensitivity map [20] assigns high attribution to the same features with a lower accuracy but assigns high attributions to nondiscriminative vegetation areas between the discriminative patterns. The attribute map produced by the DeepLIFT [6] does not assign comparable attribution for the wetlands in the middle of the image but can detect them near its borders. Gradient-weighted class activation mapping (Grad-CAM) [5] applied to the last convolution layer assigns high attribution only to the image's top-middle part with low vegetation. Integrating our approach with domain knowledge enables our approach to explain the designating patterns contributing to the concept of wilderness. Attribution maps shown in Figs. 4 and 5 prioritize bare lands, wetlands, and mountain peaks' glaciers, aligning with naturalness characteristics in protected areas in Fennoscandia [21]. This integration provides a more comprehensive explanation of the classification network's decision-making in the pattern learning phase. In addition, intersection over union (IoU) results shown

TABLE I  
AVERAGE IOU FOR THE PIXELS WITH HIGH ATTRIBUTIONS  
WITH WETLANDS AND BARE LANDS CLASSES

Method →	DeepLIFT	OSM	GraddCAM	Ours
IoU %	81.2	69.1	53.3	93.2

in Table I for high attribution pixels, including naturalness patterns, such as bare lands and wetlands, further validate the effectiveness of our integrated approach.

A combination of wetlands and bare lands have a unique combination of water, soil, and vegetation, given the fact that they are characterized by the presence of waterlogged soils, which are saturated or flooded for most of the year, which hinders the accessibility to the area and minimizes anthropogenic impact. These patterns are recognized in domain knowledge as strongly associated with the naturalness of protected areas in Fennoscandia [21]. We can conclude that our method, combined with domain knowledge, produces complete and valid explanations that align with the expert's opinions.

Although our approach introduces two challenges, namely, the high computational cost and the possibility of generating artifacts in the generated images [22], our approach provides valid and complete explanations by consistently assigning high attributions to characterize geocological patterns, such as wetlands.

#### IV. CONCLUSION

We proposed a novel, explainable machine learning approach to suppress and enhance image patterns and produce attribution maps incorporated with domain knowledge to produce valid and complete explanations. We utilize the concept of activation maximization and integrate it into a Cycle GAN's objective function to maximize the characterizing patterns of a specific class. We utilized our approach to analyze patterns in satellite imagery to better understand natural protected areas. We could show that our approach is more suitable to explain the patterns corresponding to the naturalness integrity of the protected natural areas in Fennoscandia. Our proposed method can produce complete, valid explanations of the designating geocological patterns that uphold the ecological authenticity of the protected areas more than other methods. Since our approach can be extended to multiclass classifications and various applications, we consider it a promising direction in the Earth sciences.

#### REFERENCES

- [1] S. Sarkar, "Wilderness preservation and biodiversity conservation—Keeping divergent goals distinct," *BioScience*, vol. 49, no. 5, pp. 405–412, May 1999.
- [2] B. Ekim, T. T. Stomberg, R. Roscher, and M. Schmitt, "MapInWild: A remote sensing dataset to address the question of what makes nature wild [software and data sets]," *IEEE Geosci. Remote Sens. Mag.*, vol. 11, no. 1, pp. 103–114, Mar. 2023.
- [3] T. T. Stomberg, T. Stone, J. Leonhardt, I. Weber, and R. Roscher, "Exploring wilderness characteristics using explainable machine learning in satellite imagery," 2022, *arXiv:2203.00379*.
- [4] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Computer Vision—ECCV 2014 (Lecture Notes in Computer Science)*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. Cham, Switzerland: Springer, 2014, pp. 818–833.
- [5] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, 2020.
- [6] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," 2017, *arXiv:1704.02685*.
- [7] A. Bennot, G. Franchi, J. D. Ser, R. Chatila, and N. Diaz-Rodriguez, "Greybox XAI: A neural-symbolic learning framework to produce interpretable predictions for image classification," 2022, *arXiv:2209.14974*.
- [8] P. Linardatos, V. Papastefanopoulos, and S. Kotsiantis, "Explainable AI: A review of machine learning interpretability methods," *Entropy*, vol. 23, p. 18, Dec. 2020.
- [9] A. Mahendran and A. Vedaldi, "Visualizing deep convolutional neural networks using natural pre-images," *Int. J. Comput. Vis.*, vol. 120, pp. 233–255, Dec. 2016.
- [10] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," 2017, *arXiv:1703.10593*.
- [11] A. Katzmann, O. Taubmann, S. Ahmad, A. Mühlberg, M. Sühling, and H.-M. Groß, "Explaining clinical decision support systems in medical imaging using cycle-consistent activation maximization," *Neurocomputing*, vol. 458, pp. 141–156, Oct. 2021.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015, *arXiv:1512.03385*.
- [13] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune, "Synthesizing the preferred inputs for neurons in neural networks via deep generator networks," in *Advances in Neural Information Processing Systems*, vol. 29. Red Hook, NY, USA: Curran Associates, 2016.
- [14] *Protected Planet: The World Database on Protected Areas (WDPA)*, I. UNEP-WCMC-IUCN, Cambridge, U.K., Jan. 2018.
- [15] *CORINE Land Cover 2018, Europe, 6-Yearly—Version 2020\_20u1*, European Environment Agency, Copenhagen, Denmark, 2019.
- [16] L. N. Smith and N. Topin, "Super-convergence: Very fast training of neural networks using large learning rates," *Proc. SPIE*, vol. 11006, pp. 369–386, May 2019.
- [17] S. Yun, D. Han, S. Joon Oh, S. Chun, J. Choe, and Y. Yoo, "CutMix: Regularization strategy to train strong classifiers with localizable features," 2019, *arXiv:1905.04899*.
- [18] U. Demir and G. Unal, "Patch-based image inpainting with generative adversarial networks," 2018, *arXiv:1803.07422*.
- [19] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.
- [20] T. Uchiyama, N. Sogi, S. Iizuka, K. Niinuma, and K. Fukui, "Adaptive occlusion sensitivity analysis for visually explaining video recognition networks," 2022, *arXiv:2207.12859*.
- [21] U. Gunnarsson, M. Löfroth, S. Sandring, and U. Gunnarsson, "The Swedish wetland survey: Compiled excerpts from the national final report," Naturvrdverket, Stockholm: Swedish Environmental Protection Agency, Naturvrdverket, Stockholm, Tech. Rep., 6618, 2014.
- [22] A. Aggarwal, M. Mittal, and G. Battineni, "Generative adversarial network: An overview of theory and applications," *Int. J. Inf. Manage. Data Insights*, vol. 1, no. 1, Apr. 2021, Art. no. 100004.